

IPG POWER GRID OVERVIEW

Thomas Hinke
NASA Advanced Supercomputing (NAS) Division
NASA Ames Research Center
Moffett Field, CA

IPG POWER GRID OVERVIEW AND ACKNOWLEDGMENT

This presentation will provide a brief overview of the Information Power Grid.

I would like to acknowledge that many of the slides used in this presentation are based on a set of slides prepared by Tony Lisotta, for a grid tutorial that he recently presented at Global Grid Forum 7 in Tokyo.

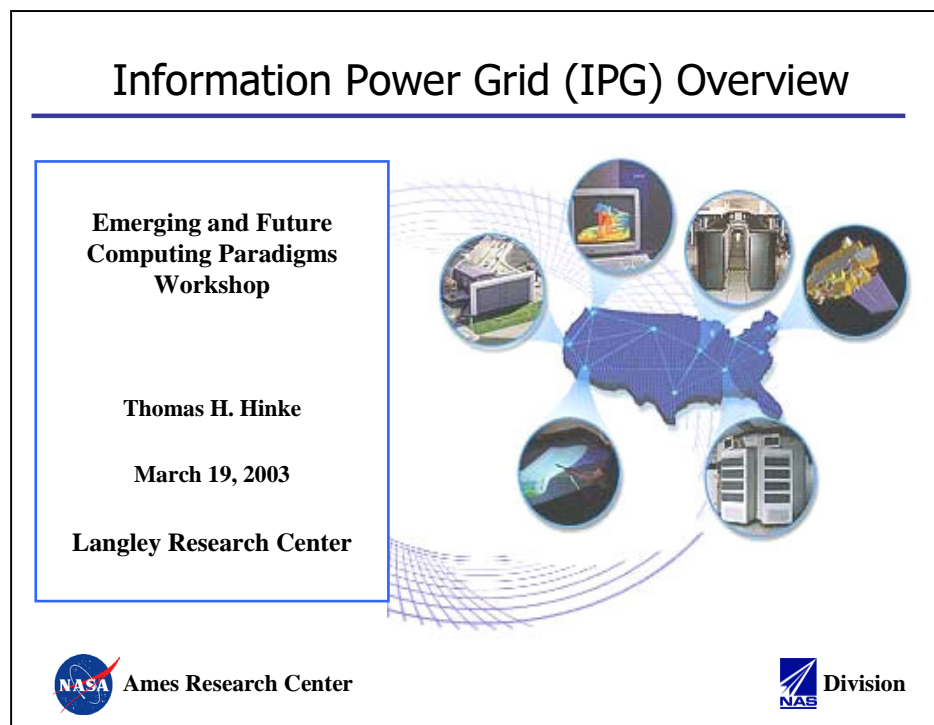


Figure 1

OUTLINE

This presentation will describe what is meant by grids and then cover the current state of the IPG. This will include an overview of the middleware that is key to the operation of the grid. The presentation will then describe some of the future directions that are planned for the IPG. Finally the presentation will conclude with a brief overview of the Global Grid Forum, which is a key activity that will contribute to the successful availability of grid components.

Outline

- What are Grids?
- Current State of Information Power Grid (IPG)
- Overview of IPG Middleware
- Future Directions
- Global Grid Forum



 Ames Research Center Division

Figure 2

WHAT DO GRIDS DO?

Grid software is middleware that sits on top of the network and the connected resources such as computers, storage and instruments. The grid software can provide an infrastructure on which to build collaborative environments that are large and distributed. They provide for security and provide the means to easily integrate distributed resources in a cost-effective manner.

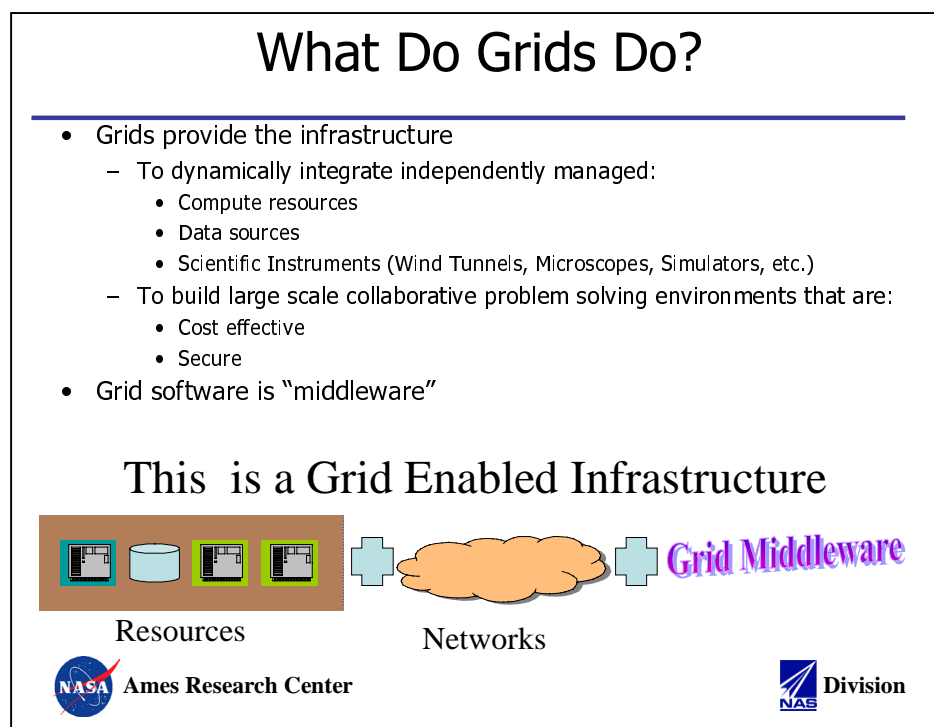


Figure 3

WHY USE GRIDS?

The goal of grids is to provide software that makes it easy for users to use distributed resources, such as distributed computers, storage or even instruments. The grid is actually a set of tools that permits these distributed resources to be easily accessed -- as if they were on the local system. These tools can also be used to develop distributed applications. They help the distributed application developer to focus on his applications, with the grid providing the software to handle the distributed access.

Why Grids?

For NASA and the general community today Grid middleware:

- Provides tools to access/use data sources (databases, instruments, ...)
- Provides tools to access computing (unique and generic)
- Is an enabler of large scale collaboration
 - Dynamically responding to needs is a key selling point of a grid.
 - Independent resources can be joined as appropriate to solve a problem.
- Provides tools for development of application-oriented frameworks
- Provides value added service to the NASA user base for utilizing resources on the grid in new and more efficient ways



 **Ames Research Center** **Division**

Figure 4

WHAT CHARACTERISTICS ARE NORMALLY FOUND IN A GRID

- Security is a fundamental aspect of a grid, with most grids basing their security on public key technology, which it used to protect at least the authentication information as it flows between the various sites on the grid. The IPG uses the Grid Security Infrastructure (GSI), based on the Globus toolkit, for its security.
- Using GSI, grids can support single sign-on, which means that after a user signs on one grid resource for a session, he is able to use other grid resources, on which he has an account, without any further identification or authentication required.
- Grids also provide a grid information service (GIS), that provides a single mechanism by which users can discover grid resources and associated information about the resource.
- Grids are designed to be scalable to a large number of resources.
- Finally, grids are designed to provide access to resources that may be under the control of different administrative groups. They are not designed to have centralized control.

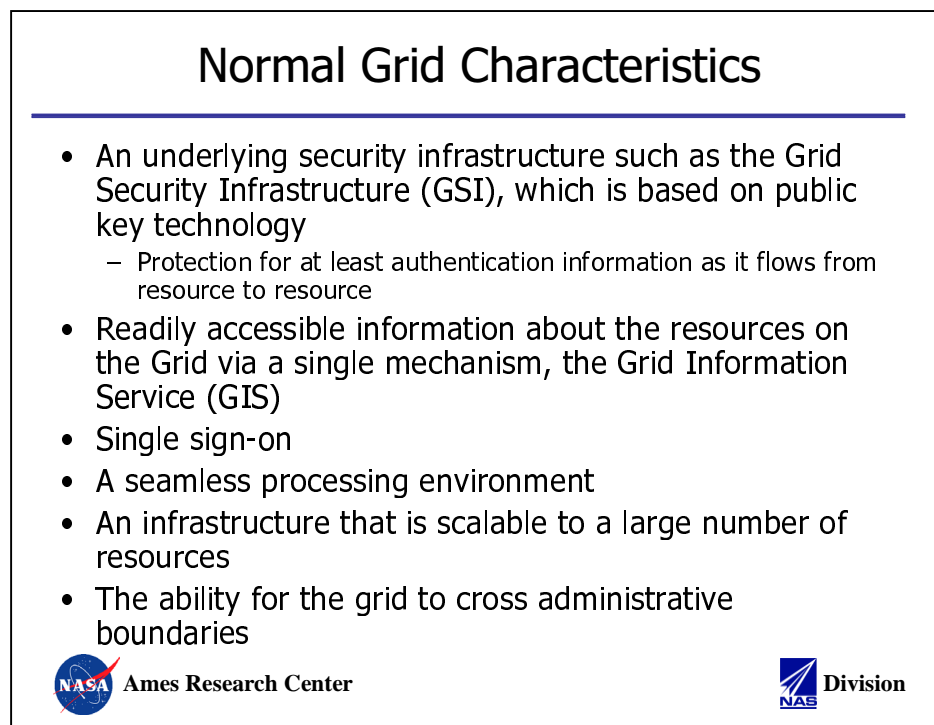


Figure 5

DISTRIBUTED SYSTEMS BEFORE THE GRID

Before the development of the grid, people still developed distributed systems. Under these pre-grid distributed systems, a user was responsible for dealing with all of the complexities of the distributed environment.

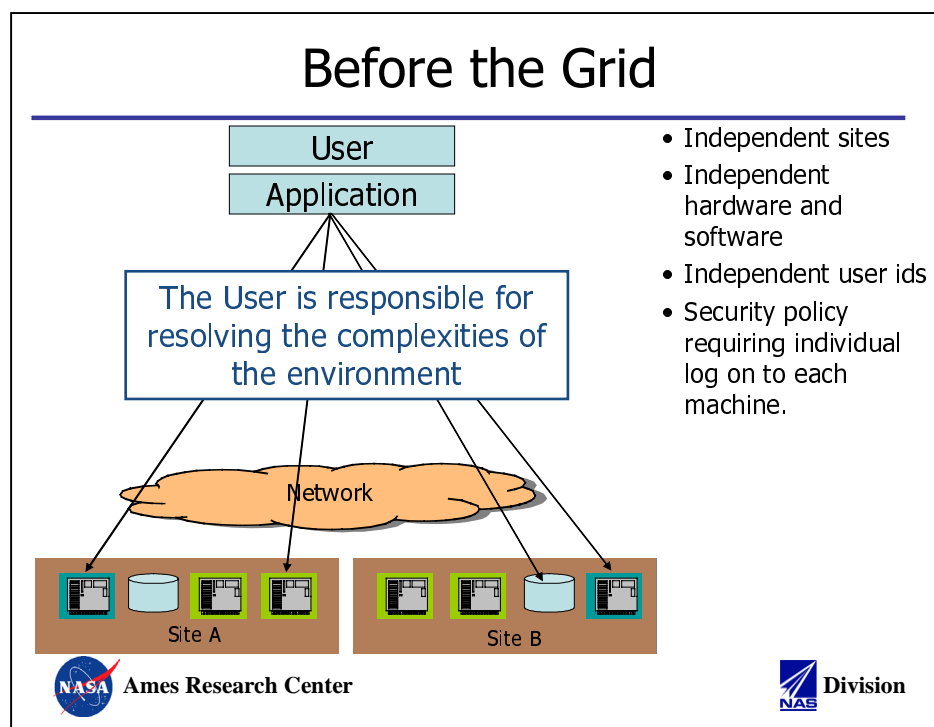


Figure 6

DISTRIBUTED SYSTEMS USING TODAY'S GRID

The grid provides the middleware that ties distributed resources into a seamless environment. Using the grid, a user can make a request to the grid Information Service for information about the location and characteristics of grid resources such as processing and storage resource or instruments. With this information, the user can then launch an application that accesses the desired distributed resources through the grid middleware.

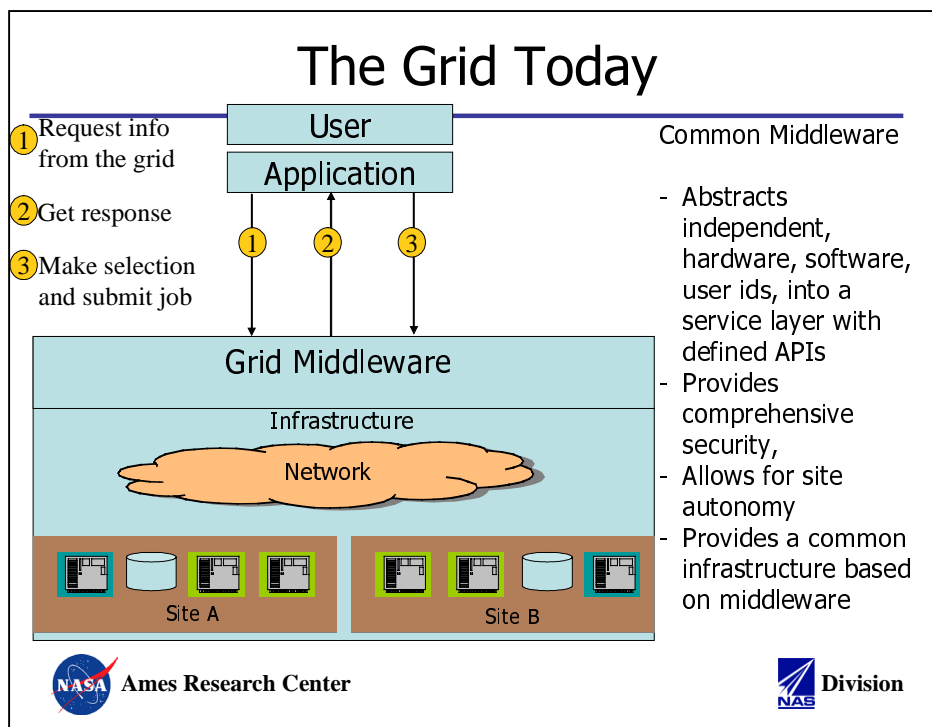


Figure 7

DISTRIBUTED SYSTEMS USING TODAY'S GRID

The key to the grid is that the underlying grid resources are abstracted into application programmer interfaces that simplify the development of distributed applications. While this is a significant step forward, this layer does not have much intelligence, which will define the next stage of grid development.

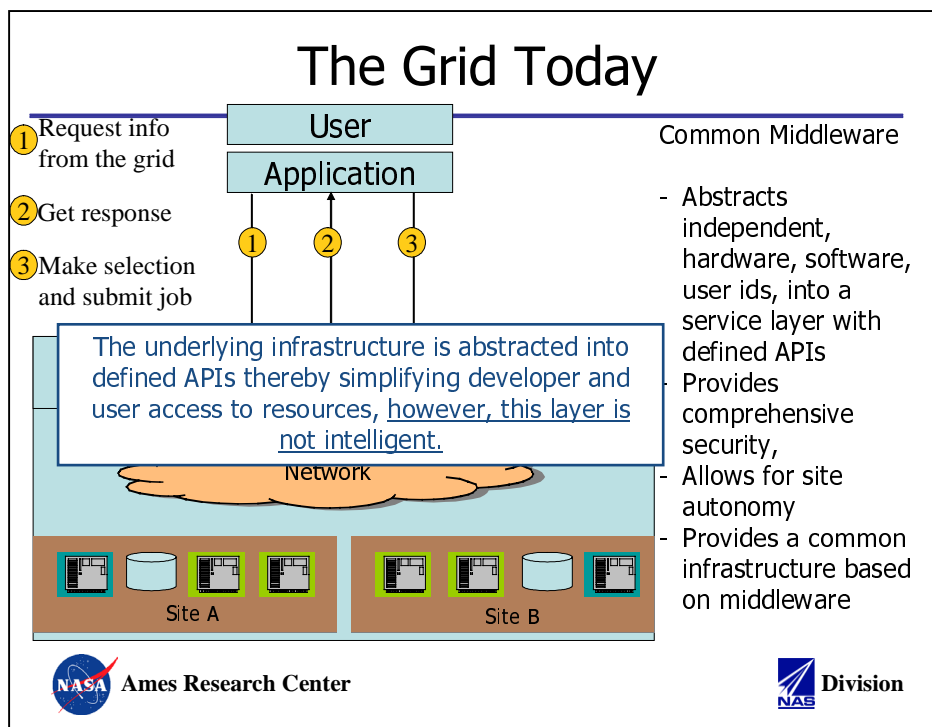


Figure 8

THE NEAR FUTURE GRID WILL HAVE INTELLIGENCE

The grid for the near future will have intelligent, customizable middleware that will sit between the current grid middleware and the application. This intelligent layer will perform brokering (the automatic selection of resources) and will provide information tailored to the specific needs of the user or application.

Under the current grid, a user must have an account on each resource that is used, thus preserving local autonomy. Under the near future grid, if a local system agrees, the grid will then take responsibility for granting grid user's access to these resources, where the user has not pre-established an account.

Another key capability that will soon be available is the ability to field grid-enabled web services, that provide a standard API that can be accessed from applications, application-specific portals or command-line functions.

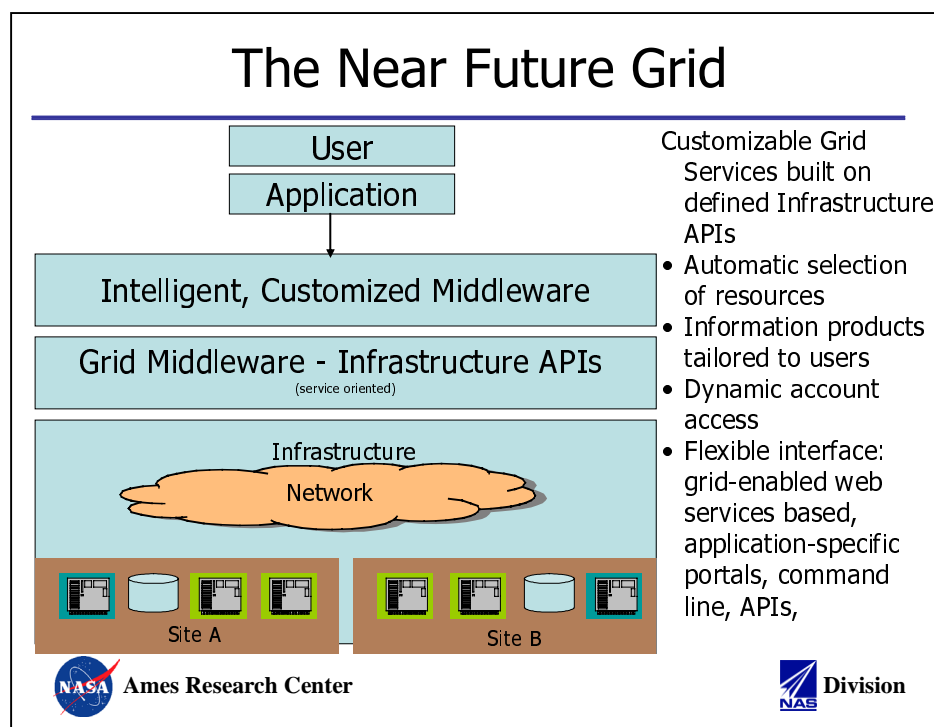


Figure 9

THE NEAR FUTURE GRID WILL HAVE INTELLIGENCE

With this more intelligent grid, the users and application developers will be able to focus more on the science and engineering applications and not on the distributed systems management aspects of their systems.

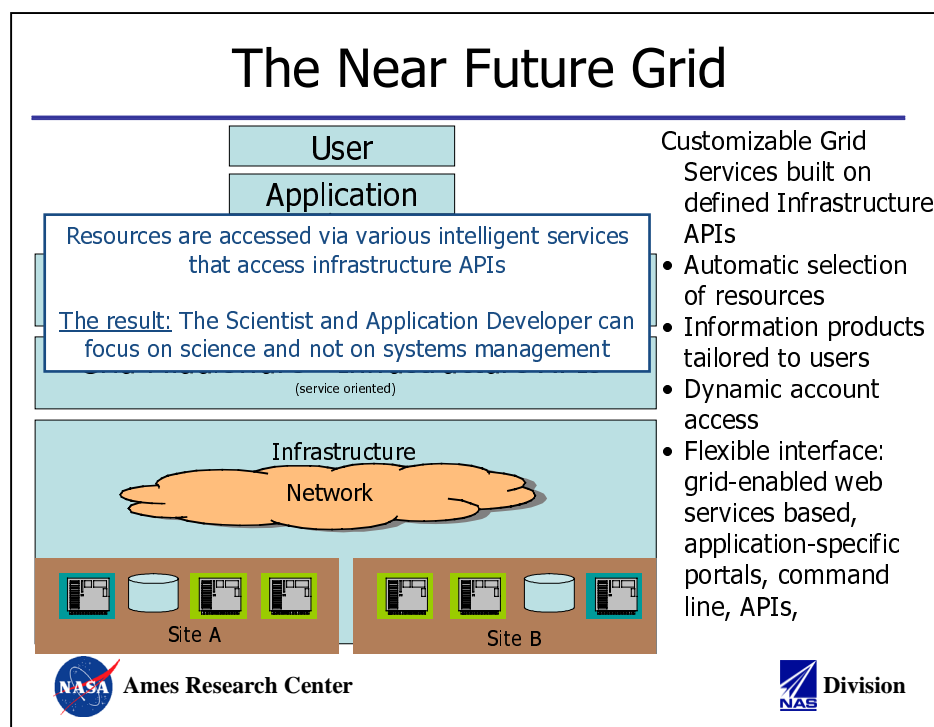


Figure 10

HOW THE USER AND APPLICATION DEVELOPERS SEE A GRID

A grid is really just a set of tools that can be accessed through application programmer interfaces or command line functions. These tools will be augmented with services that will be structured as grid-enabled web services, which are re-usable such that one or more of these can be combined to make a more complex services.

Once a user has authenticated to the grid, he can use any of the various services that are shown on the slide as if these were part of his local machine. He does not have to re-authenticate to use any of these, with the grid handling the requirement to pass identification and authentication information among the resources that are used.

How the User and Application Developers See a Grid

- A set of grid functions that are available as
 - Application programmer interfaces (APIs)
 - Command-line functions
 - Grid-enabled web services
- After authentication, grid functions can be used to
 - Spawn jobs on different processors with a single command
 - Access data on remote systems
 - Move data from one processor to another
 - Support the communication between programs executing on different processors
 - Discover the properties of computational resources available on the grid using the grid information service
 - Use a broker to select the best place for a job to run and then negotiate the reservation and execution (coming soon).



 Ames Research Center Division

Figure 11

OUTLINE

In the next section we will look at the current state of the IPG.

Outline

- What are Grids?
- Current State of Information Power Grid (IPG)
- Overview of IPG Middleware
- Future Directions
- Global Grid Forum



 Ames Research Center Division

Figure 12

IPG LOCATIONS

The IPG currently has resources located at the five NASA Centers shown on the map.

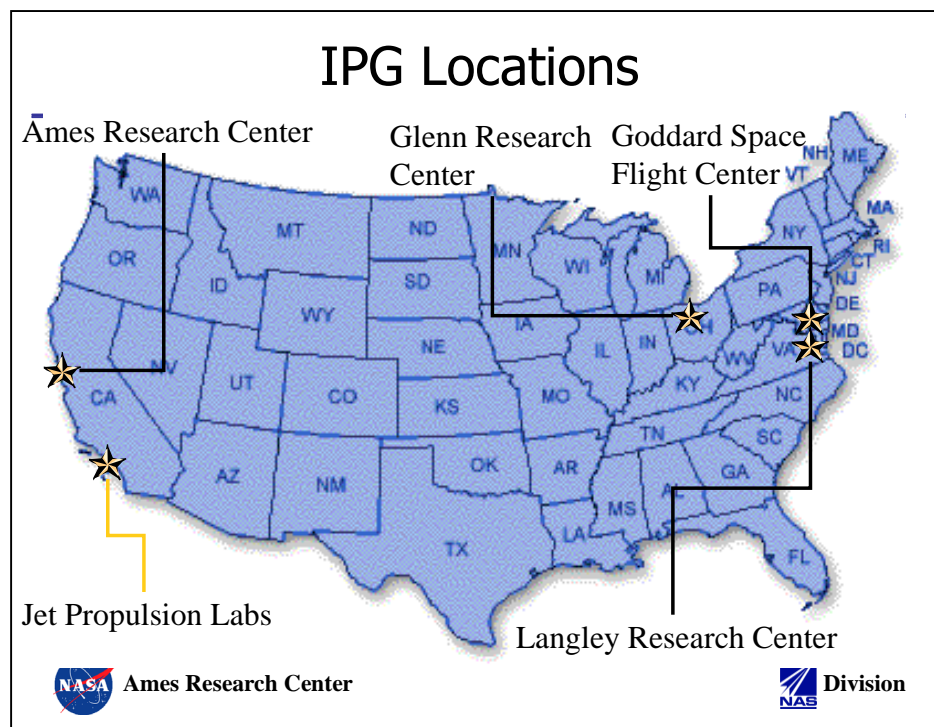


Figure 13

IPG RESOURCES

The IPG currently has the computational resources shown.

IPG Resources

- Server Nodes
 - 1024 CPU, single system image SGI, **Ames**
 - 512 CPU SGI O2K, **Ames**
 - 128 CPU Linux Cluster, **Glenn**
 - 124 CPU SGI O2K, **Ames**
 - 64 CPU SGI O2K, **Ames**
 - 24 CPU SGI O2K, **Glenn**
 - 16 CPU SGI O2K, **Langley**
 - 16 CPU SGI O2K, **Ames**
 - 8 CPU SGI, O3K, **Langley**
 - 4 CPU SGI O2K, **Langley**
- Client Nodes
 - 16 CPU SGI O300, JPL
 - 8 CPU SGI O300, Goddard
- Wide area network interconnects of at least 100 Mbit/s



 Ames Research Center Division


Figure 14

OUTLINE

The next section will delve more deeply into the nature of the IPG middleware.

Outline

- What are Grids?
- Current State of Information Power Grid (IPG)
- Overview of IPG Middleware
- Future Directions
- Global Grid Forum

 NASA Ames Research Center


 NAE Division

Figure 15

IPG IS BUILD ON GLOBUS TOOLKIT 2

The IPG, as are most of the grids in the world, is built on Globus Toolkit 2 (GT2). The Grid Security Infrastructure (GSI) is based on X509 certificates, secure socket layer (SSL) and Transfer Layer Security (TLS). This supports a GSI-enabled Secure Shell (SSH) and GridFTP (a high performance GSI version of FTP).

The Grid Information Services is based on LDAP (lightweight Directory Access Protocol) which supports the Monitoring and Discovery Service (MDS), which provides a directory of grid resources and attributes.

Finally, the remote execution of jobs is supported by the Globus Resource Allocation Manager (GRAM), which provides an interface to various batch schedulers (e.g., PBS and LSF), as well as systems that permits users to directly execute jobs via fork. It permits the launching of remote jobs.

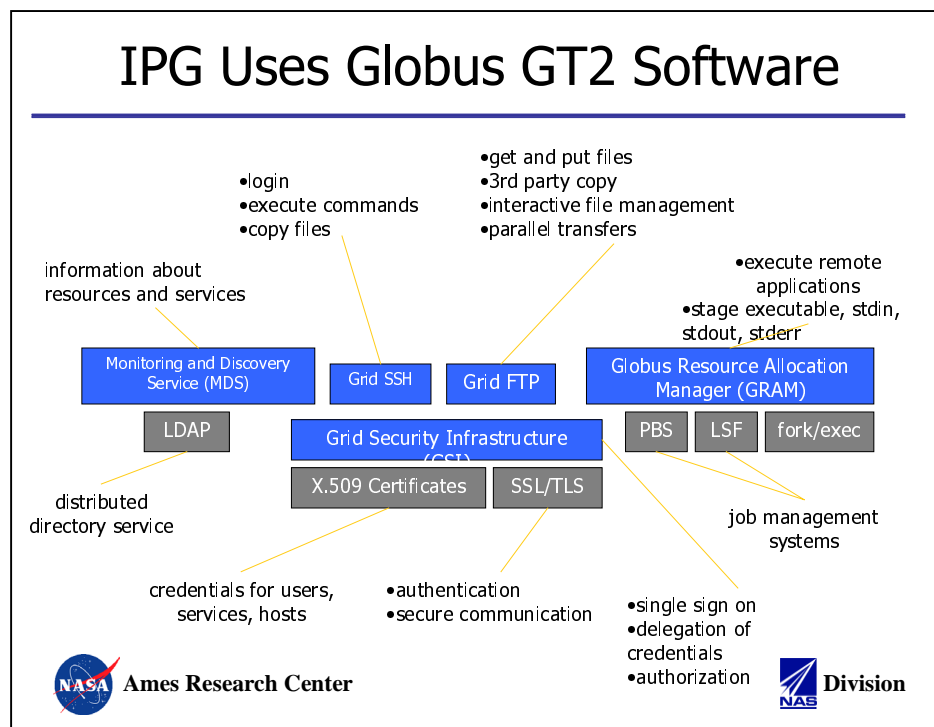


Figure 16

IPG/GLOBUS DEPLOYMENT ARCHITECTURE

To support the grid information service of a deployed grid, a Grid Resource Information Service (GRIS) captures local information from each resource and forwards this to a Grid Index Information Service (GIIS), that provides a single source for information about a particular grid.

Users, applications or web portals can use Globus client services to access any of the grid tools and services.

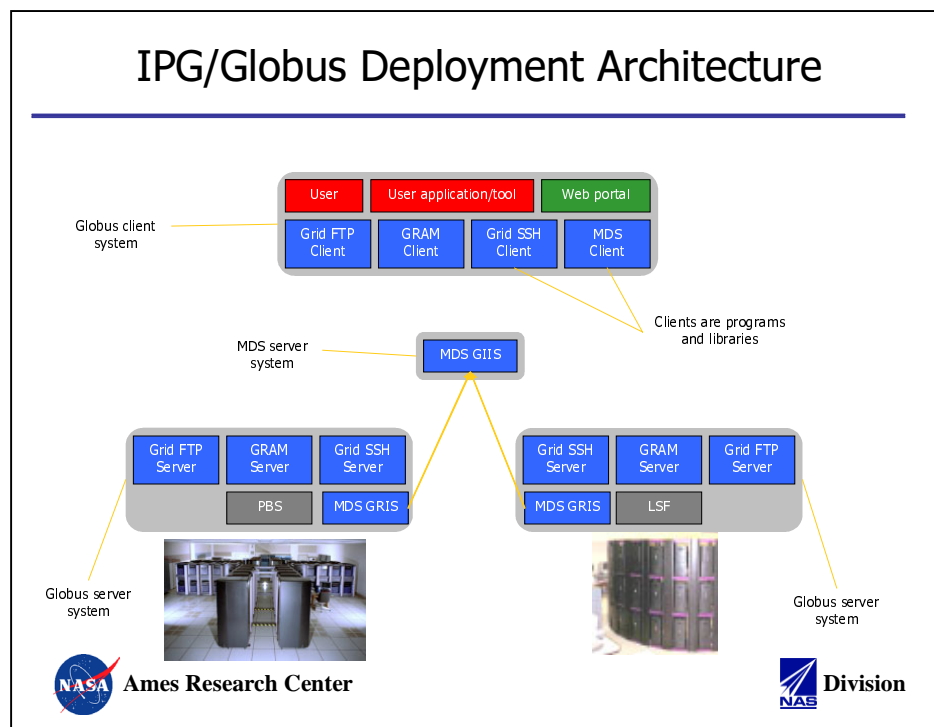


Figure 17


ADDITIONAL SERVICE UNDER DEVELOPMENT BY THE IPG PROJECT

To provide the added intelligence needed to facilitate the development of grid applications and the use of the grid by users, the IPG project is developing a Job Manager to manage the reliable execution of a job on the grid. The Job Manager will stage the necessary files needed by the application, monitor the progression of the work and then post-stage the results, cleaning up any files that may remain from the execution.

The Job Manager is supported by the Resource Broker that provides the user with suggestions about where to run his application, based on supplied information about the application.

Additional IPG Services

- Job Manager
 - Reliably execute a job
 - Set of files to pre-stage
 - Executable to run
 - Including directory, environment variables
 - Set of files to post-stage
- Resource Broker
 - Provide suggestions on where to run a job
 - Input
 - Which hosts and operating systems are acceptable
 - How to create a Job Manager Job for a selected host
 - Selection made using host and OS constraints and host load
 - Interactive system: # free CPUs
 - Batch system: Amount of work in queue / # CPUs
 - Output
 - Ordered list of Job Manager Jobs (suggested systems)

 **Ames Research Center**


 **Division**

Figure 18

ROLE OF ADDITIONAL IPG SERVICES

Applications will be able to consult the broker for suggestions as to the best grid resources to use, given the current workload on each of these resources. This information will then be used to run the application on the suggested resources, using the job manager to stage necessary files and monitor the progress of the work and then post stage any files at the end of the work.

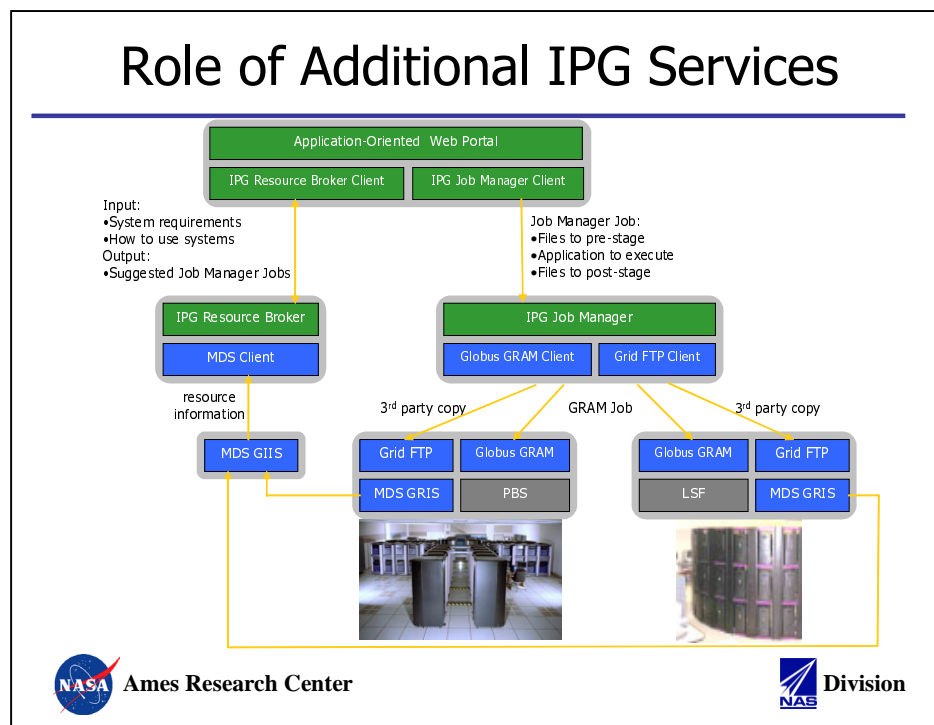



Figure 19

OUTLINE

Next we will briefly look at future directions.

Outline

- What are Grids?
- Current State of Information Power Grid (IPG)
- Overview of IPG Middleware
- Future Directions
- Global Grid Forum

 NASA Ames Research Center


 NAE Division

Figure 20

OPEN GRID SERVICES ARCHITECTURE (OGSA)

The Open Grid Services Architecture is the grid community's adoption of the web services work (which other than the name has little to do with the web) as a way of delivering services. Grid-enabled web services provide a standard Web Services Description Language (WSDL) description of the service and a specified protocol, which for now is SOAP, for accessing these services. Grid-enabled web services provide a self-describing way to offer services that can be included as components of other grid-enabled web service.

Standards are under development by the Global Grid Forum to specify the interfaces and the nature of the service-management capabilities (creation, destruction, lifetime) that are to be associated with each service.

One of the key contributions that grid-enabled web services offer over web services is that they will be built to use grid security, such as the Grid Security Infrastructure.

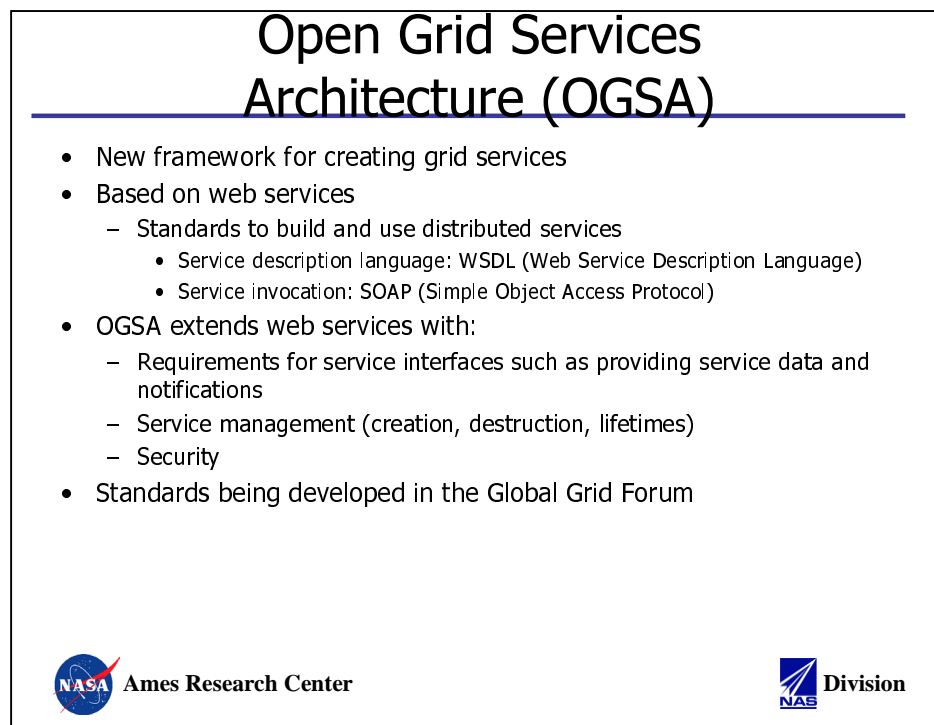


Figure 21

GLOBUS TOOLKIT VERSION 3 (GT3)

A key first application of OGSA will be the next version of the Globus Toolkit, which is called Globus Toolkit Version 3 (GT3). The various grid services offered by the Globus Toolkit will be offered as grid-enabled web services.

GT3 and OGSA will revolutionize how services are offered on the grid, since it will make it easy to include existing services in more complex, application-specific services.

The IPG will transition to GT3 as soon as it is stable and in a way that minimizes any impact to existing users.

Globus Toolkit Version 3 (GT3)

- Large change from GT2 to GT3
 - New implementation
 - Java-based instead of C-based
 - GT3 based on OGSA
- GT3 will provide equivalent services to GT2
- Alpha version of GT3 currently available
- GT3 and OGSA will revolutionize
 - how services are provided on the grid and
 - how grid applications are developed
- IPG will transition to GT3 soon as it is proven stable, while minimizing the effect on existing IPG users.
- Transition should have minimal impact on IPG users
 - Globus will maintain many of the existing programs
- IPG Services will follow OGSA



 **Ames Research Center** **Division**

Figure 22

FOCUS ON IPG HANDLING OF DATA

As the IPG completes its work on the resource management and utilization phase of the grid services, it will focus on the data handling aspects of the grid. This is a critical function for NASA because of the large volume of distributed data that is found in the various NASA archives, such as those associated with Earth science.

This new focus will look at providing access to NASA archives, using such existing grid-enabled systems as the Storage Resource Broker, developed at the San Diego Supercomputing Center. Of particular interest will be providing access to data stored on both tertiary storage (mass storage systems) and data stored on disk-resident data pools.

This effort will build on the considerable amount of work that has been performed on data grids by the international grid community.

Focus on IPG Handling of Data

- Goal: Intelligently manage data in a grid
- NASA data is inherently distributed e.g., various Earth science archives, including the one at LaRC
- Important focus of IPG
- Access to files
 - Initial use of grid-enabled Storage Resource Broker
 - Data staging and replica management building on grid community research
 - Need grid support for file metadata
- NASA data can be on
 - Disk-resident data pools
 - Tertiary storage data archives
- Will build on considerable data grid work from the international grid community



Ames Research Center



Division


Figure 23

OUTLINE

The last section will focus on the Global Grid Forum.

Outline

- What are Grids?
- Current State of Information Power Grid (IPG)
- Overview of IPG Middleware
- Future Directions
- Global Grid Forum

 NASA Ames Research Center


 NAE Division

Figure 24

GLOBAL GRID FORUM BACKGROUND

The Global Grid Forum is an international group that mirrors for grids what the Internet Engineering Task Force (IETF) has done for the network through its standards work. It was formed in 2001 as a combination of similar grid work in the North America and Europe and now encompasses the Asia/Pacific grid work as well. It meets three times a year in different parts of the world.

Global Grid Forum Background

- Began in 2001 as merger of previous regional grid forums.
- Now includes grid technical communities in North America, Europe and Asia Pacific
- Meets three times per year, alternating between North America and Europe and Asia/Pacific
- Modeled after IETF (Internet Engineering Task Force), which sets Internet standards.
- GGF7 was just held in Tokyo, Japan with over 700 attendees
- GGF8 will be held in Seattle, WA in June 25-27, 2003



Ames Research Center



Division

Figure 25

GLOBAL GRID FORUM PURPOSE AND ORGANIZATION

The main purpose of the Global Grid Forum is to provide an international grid organization that can support the fair and representative development, review, approval and release of both best practices and standards for the grid.

It is organized into two types of groups. The Working Groups are of limited duration and are focused on the goal of producing some specific best practice document or standard. Currently there are 24 Working Groups.

The Research Groups are organized to address grid issues that are not yet ready for a best practice document or a standard. Currently there are 20 research groups.





Figure 26

GGF WORKING GROUPS

The slide lists the current GGF Working Groups. Details about each of these groups and the current set of documents and standards on which they are working can be found on the GGF web site at www.ggf.org.

GGF Working Groups

•Grid Checkpoint Recovery	•Discovery and Monitoring Event Description
•New Productivity Initiative	•Network Measurement
•Open Grid Services Architecture	•Grid Information Retrieval
•Open Grid Services Interface	•Previous activities of the Peer to Peer Working Group
•Open Source Software	•Distributed Resource Management Application API
•Data Access & Integration Services	•Grid Economic Services Architecture
•GridFTP	•Grid Resource Allocation Agreement Protocol
•Authorization Frameworks and Mechanisms	•OGSA Resource Usage Service
•Certificate Authority Ops	•Scheduling Attributes
•Grid Certificate Policy	•Scheduling Dictionary
•Grid Security Infrastructure	•Usage Record
•Open Grid Service Architecture Security	
•CIM based Grid Schema	

 Ames Research Center


 Division


Figure 27

GGF RESEARCH GROUPS

The slide lists the current GGF Research Groups. Details about each of these groups can be found on the GGF web site at www.ggf.org.

GGF Research Groups

• Advanced Collaborative Environments	• Data Replication
• Advanced Programming Models	• Data Transport
• Applications and Test Beds	• Grid Benchmarking
• Grid Computing Environments	• Relational Grid Information Services
• Grid User Services	• Appliance Aggregation (
• Life Sciences Grid	• OGSA-P2P-Security
• Production Grid Management	• Grid High-Performance Networking
• Accounting Models	• Persistent Archives
• Grid Protocol Architecture	• Site Authentication, Authorization, and Accounting Requirements
• Semantic Grid	
• Service Management Frameworks	

 **Ames Research Center**


 **Division**

Figure 28

WHY IS THE GLOBAL GRID FORUM IMPORTANT

The primary reason that the GGF is important is that it will result in grid standards and grid standards will encourage commercial companies to make grid products that satisfy these standards. Standard based products should be more marketable than products that do not satisfy standards.

In addition the GGF provides an arena for various application-specific requirements to be injected into the international grid community. Currently there are a number of application-specific research groups at GGF that may, as the need is found, develop application-specific standards or influence other standards work to address needs unique to a particular application area.

Why is the Global Grid Forum Important

- It will result in grid standards
 - It will encourage commercial products since there will be standards which the products can meet
 - Products that meet accepted standards should be more marketable
- It provides a forum to get application-specific requirements injected into the grid development efforts



 Ames Research Center Division

Figure 29